**Robert B. Ross**
9700 S. Cass Ave.
Building 240
Argonne, IL 60439-4844

email: robertbross@acm.org
email: rross@mcs.anl.gov

**I am a computer scientist and engineer fascinated with the design, implementation, and deployment of resilient and high performance distributed systems.** My focus area for nearly two decades has been data and communication system software for high performance computing. Over that time I have participated in and lead the development of a number of algorithms and open source software packages for scientific computing.

## Professional Preparation

| | | |
|---|---|---|
| Clemson University, Clemson, SC | Computer Engineering, Focus: Computer Architecture | Ph.D., 2000 |
| Clemson University, Clemson, SC | Computer Engineering, Minor in Math Science | B.S., 1994 |

## Appointments

| | | |
|---|---|---|
| 2014–present | Senior Computer Scientist | Mathematics and Computer Science Division, Argonne National Laboratory |
| 2004–2014 | Computer Scientist | Mathematics and Computer Science Division, Argonne National Laboratory |
| 2012–present | Senior Fellow | Northwestern–Argonne Institute for Science and Engineering, Northwestern University and Argonne National Laboratory |
| 2011–present | Senior Fellow | Computation Institute, The University of Chicago and Argonne National Laboratory |
| 2004–present | Adjunct Assistant Professor | Department of Electrical and Computer Engineering, Clemson University |
| 2004–2011 | Fellow | Computation Institute, The University of Chicago and Argonne National Laboratory |
| 2002–2004 | Assistant Computer Scientist | Mathematics and Computer Science Division, Argonne National Laboratory |

## Current Activities

**Darshan I/O Characterization Tool**
Lightweight tool for observing the I/O behavior of HPC applications at the largest scales.

**CODES Simulation Framework**
Parallel discrete event simulation framework, building on RPI's ROSS system, that enables high fidelity simulation of complex distributed storage systems.

**Triton Distributed Storage System**
Prototype, under development, of an object based, distributed storage system for use in computational and data intensive science.

**SciDAC Scalable Data Management, Analysis, and Visualization Institute (SDAV)**
Deputy Director, working with computational scientists and other experts in scientific data management and analysis to support data management solutions for DOE computational science applications.

**DOE Office of Advanced Scientific Computing Research Extreme Scale Systems Research**
ASCR lead for Data Management for extreme scale computing, providing guidance and coordinating research activities in this area along with NNSA co-lead.

**Argonne Math and Computer Science Division Data Strategy**
Strategic Area Lead, working with computer scientists and applied mathematicians to develop division strategy for Big Data and data-driven science research and development in support of the DOE mission.

## Honors and Awards

| | |
|---|---|
| 2011 | LSAP 2011 Best Paper award for "Visual Analysis of I/O System Behavior for High End Computing" |
| 2011 | MSST 2011 Best Paper award for "Understanding and Improving Computational Science Storage Access through Continuous Characterization" |
| 2009 | EuroPVM/MPI 2009 Outstanding Paper award for "Processing MPI Datatypes Outside MPI" |
| 2008 | Clemson University College of Engineering and Science Outstanding Young Alumni |
| 2005 | R&D 100 Award Winner for MPICH2 |
| 2004 | Presidential Early Career Award for Scientists and Engineers<br>Dept. of Energy Office of Science Early Career Scientist and Engineer Award |
| 2000–2002 | Argonne National Laboratory Enrico Fermi Scholar |
| 1999 | USENIX 2000 Best Paper award for "PVFS: A Parallel File System for Linux Clusters" |
| 1996–1999 | NASA Graduate Student Research Program Fellow |

## Professional Activities

IDC Technical Computing Advisory Panel 2014

IPDPS PC 2014, reviewer 2002, 2007

IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid) PC 2015

Workshop on Many-Task Computing on Clouds, Grids, and Supercomputers (MTAGS) PC 2014

IEEE Conference on Massive Data Storage (MSST) PC 2013-2014

Parallel Data Storage Workshop (PDSW) General Chair 2013, Steering Committee 2013-2014, PC Chair 2012, PC 2009–2014

IEEE Transactions on Computers reviewer 2013

NSF PRObE steering committee 2011

International Workshop on Runtime and Operating Systems for Supercomputers (ROSS) PC 2011–2014

IEEE Cluster steering committee 2010–2013; PC 2003, 2006, 2009-2011, 2015; reviewer 2002

International Workshop on Data Intensive Computing in the Cloud (DataCloud) PC 2013

Workshop on Big Data Management in Clouds PC 2013

USENIX Conference on File and Storage Technologies (FAST) PC 2011; reviewer 2005

IEEE Symposium on Large Scale Data Analysis and Visualization (LDAV) PC 2011–2012

EuroMPI PC 2010–2012; session chair 2006

HotStorage PC 2012

NSF review panelist 2006, 2007, 2010, 2012

NSF external reviewer 2011

International Conference on Parallel Architectures and Compilation Techniques (PACT) reviewer 2012

SciDAC Center for Scalable Application Development Software (CScADS) Workshop organizer 2008–2011

International Conference on Distributed Computing Systems (ICDCS) PC 2010

International Conference on Parallel and Distributed Systems (ICPADS) PC 2010

SciDAC Conference PC 2006, 2009, 2010

IEEE Transactions in Parallel and Distributed Systems reviewer 2003-2010

Computational Science and Discovery editorial board 2010

Workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS) workshop chair 2009

International Journal of High Performance Computing and Applications (IJHPCA) reviewer 2009, 2011–2012

Programmable File Systems Workshop PC 2014

DOE/ASCR 2007 Visualization and Analytics Workshop co-chair

SC storage co-chair 2011; PC 2004, 2007–2009, 2012-2013; reviewer 2003

DOE Office of Science ASCR PI Meeting steering committee 2008

International Conference on Parallel Processing (ICPP) PC 2007, 2009, 2013

DOE SBIR reviewer 2007, 2009

International Journal of Computers and Their Applications (IJCA) reviewer 2006

Journal of Parallel and Distributed Computing reviewer 2003, 2006

Journal of Parallel Computing reviewer 2005

High Performance Distributed Computing Conference reviewer 2005

DOE Office of Science Early Career Principal Investigator program reviewer 2005

NASA Computing, Networking, and Information Systems R&D activities reviewer 2004

Special issue of Cluster Computing Journal reviewer 2003

DOE Office of Science Early Career program reviewer 2002

## Supervisory and Mentoring Activities

- Ph.D. thesis committee member for Philip Carns (2005), Murali Vilayannur (2005), Avery Ching (2007), Dries Kimpe (2008), Sumit Narayan (2010), Arifa Nisar (2010), Jing Fu (2012), Zhiwei Sun (2013), Misbah Mubarak (TBD), Michael Kasick (TBD), and Huong Luu (TBD)
- Software programmer supervisor 2003–present
- Postdoctoral researcher supervisor 2004–present
- Summer student supervisor 2001–present

## Book Chapters and Journal Articles


[1] J. Jenkins, X. Zou, H. Tang, D. Kimpe, R. Ross, and N. F. Samatova. RADAR: Runtime asymmetric data-access driven scientific data replication. In *Supercomputing*, pages 296–313. Springer International Publishing, 2014.

[2] T. Ilsche, J. Schuchart, J. Cope, D. Kimpe, T. Jones, A. Knupfer, K. Iskra, R. Ross, W. Nagel, and S. Poole. Optimizing I/O forwarding techniques for extreme-scale event tracing. *Cluster Computing*, pages 1–18, 2013.

[3] J. Jenkins, I. Arkatkar, S. Lakshminarasimhan, D. A. Boyuka II, E. R. Schendel, N. Shah, S. Ethier, C.-S. Chang, J. Chen, H. Kolla, et al. ALACRITY: Analytics-driven lossless data compression for rapid in-situ indexing, storing, and querying. In *Transactions on Large-Scale Data-and Knowledge-Centered Systems X*, pages 95–114. Springer Berlin Heidelberg, 2013.

[4] R. Latham and R. Ross. Parallel I/O basics. In *Earth System Modelling-Volume 4*, pages 3–12. Springer Berlin Heidelberg, 2013.

[5] N. Liu, C. Carothers, J. Cope, P. Carns, and R. Ross. Model and simulation of exascale communication networks. *Journal of Simulation*, March 2012.

[6] S. Lakshminarasimhan, N. Shah, S. Ethier, S.-H. Ku, C. Chang, S. Klasky, R. Latham, R. Ross, and N. F. Samatova. ISABELA for effective in situ compression of scientific data. *Concurrency and Computation: Practice and Experience*, 2012.

[7] R. Latham, C. Daley, W. keng Liao, K. Gao, R. Ross, A. Dubey, and A. Choudhary. A case study for scientific I/O: improving the FLASH astrophysics code. *Computational Science and Discovery*, 5(1):015001, 2012.

[8] P. Carns, K. Harms, W. Allcock, C. Bacon, S. Lang, R. Latham, and R. Ross. Understanding and improving computational science storage access through continuous characterization. *ACM Transactions on Storage*, 7(3), October 2011.

[9] R. Ross. Parallel file systems. In D. Padua, editor, *The Encyclopedia of Parallel Computing*. Springer, September 2011.

[10] F. Isaila, J. G. Blas, J. Carretero, R. Latham, and R. Ross. Design and evaluation of multiple level data staging for bluegene systems. *IEEE Transactions on Parallel and Distributed Systems*, 22(6), June 2011.

[11] R. Ross, A. Choudhary, G. Gibson, and W.-K. Liao. Parallel data storage and access. In A. Shoshani and D. Rotem, editors, *Scientific Data Management: Challenges, Technology, and Deployment*. Chapman & Hall/CRC, 2010.

[12] R. Ross, P. Carns, and D. Metheney. Parallel file systems. In Y. Chan, J. Talburt, and T. Talley, editors, *Data Engineering: Mining, Information and Intelligence*. Springer, October 2009.

[13] R. Latham, R. B. Ross, and R. Thakur. Implementing MPI-IO atomic mode and shared file pointers using mpi one-sided communication. *Int'l Journal of High Performance Computing Applications*, 21(2):132–143, Summer 2007.

[14] A. Ching, A. Choudhary, W. K. Liao, R. Ross, and W. Gropp. Evaluating structured I/O methods for parallel file systems. *International Journal of High Performance Computing and Networking*, 2:133–145, 2004.

[15] R. B. Ross and W. B. L. III. Server-side scheduling in cluster parallel I/O systems. In C. Cerin and H. Jin, editors, *Parallel I/O for Cluster Computing*, pages 157–178. Kogan Page Science, Sterling, VA, 2004.

[16] W. B. L. III and R. B. Ross. Parallel I/O and the Parallel Virtual File System. In W. Gropp, E. Lusk, and T. Sterling, editors, *Beowulf Cluster Computing with Linux, second edition*, pages 493–534. MIT Press, Cambridge, MA, 2003.

[17] W. B. L. III and R. B. Ross. PVFS: Parallel Virtual File System. In T. Sterling, editor, *Beowulf Cluster Computing with Linux*, pages 391–429. MIT Press, Cambridge, MA, 2002.

## Refereed Proceedings

[1] W. Tang, J. Jenkins, F. Meyer, R. Ross, R. Kettimuthu, L. Winkler, X. Yang, T. Lehman, and N. Desai. Data-aware resource scheduling for multicloud workflows: A fine-grained simulation approach. In *Emerging Issues in Cloud Workshop, held in conjunction with the 6th IEEE International Conference on Cloud Computing Technology and Science*, December 2014.

[2] D. Dai, Y. Chen, D. Kimpe, and R. Ross. Two-choice randomized dynamic I/O scheduler for object storage systems. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC14)*, November 2014.

[3] M. Dorier, S. Ibrahim, G. Antoniu, R. Ross, et al. Omnisc'IO: A grammar-based approach to spatial and temporal I/O patterns prediction. In *International Conference for High Performance Computing, Networking, Storage and Analysis (SC14)*, November 2014.

[4] S. Snyder, P. Carns, J. Jenkins, K. Harms, R. Ross, M. Mubarak, and C. Carothers. A case for epidemic fault detection and group membership in HPC storage systems. In *Proceedings of the 5th International Workshop on Performance Modeling, Benchmarking, and Simulation of High Performance Computer Systems (PMBS14)*, November 2014.

[5] D. Dai, Y. Chen, D. Kimpe, R. Ross, and X. Zhou. Domino: an incremental computing framework in cloud with eventual synchronization. In *Proceedings of the 23rd international symposium on High-performance parallel and distributed computing*, pages 291–294. ACM, June 2014.

[6] D. Dai, Y. Chen, D. Kimpe, and R. Ross. Provenance-based prediction scheme for object storage system in hpc. In *Cluster, Cloud and Grid Computing (CCGrid), 2014 14th IEEE/ACM International Symposium on*, pages 550–551. IEEE, May 2014.

[7] M. Dorier, G. Antoniu, R. Ross, D. Kimpe, and S. Ibrahim. CALCioM: Mitigating I/O interference in HPC systems through cross-application coordination. In *Proceedings of the International Parallel and Distributed Processing Symposium*, May 2014.

[8] M. Mubarak, C. D. Carothers, R. B. Ross, and P. Carns. A case study in using massively parallel simulation for extreme-scale torus network codesign. In *Proceedings of the 2nd ACM SIGSIM/PADS conference on Principles of advanced discrete simulation*, pages 27–38. ACM, May 2014.

[9] D. Dai, R. Ross, P. Carns, D. Kimpe, and Y. Chen. Using property graphs for rich metadata management in hpc systems. In *Proceedings of the 9th Parallel Data Storage Workshop*, pages 7–12. IEEE Press, 2014.

[10] S. Kumar, A. Saha, V. Vishwanath, P. Carns, J. A. Schmidt, G. Scorzelli, H. Kolla, R. Grout, R. Latham, R. Ross, et al. Characterization and modeling of PIDX parallel I/O for performance optimization. In *Proceedings of SC13: International Conference for High Performance Computing, Networking, Storage and Analysis*, page 67. ACM, November 2013.

[11] J. Lofstead and R. Ross. Insights for exascale IO APIs from building a petascale IO API. In *Proceedings of SC13: International Conference for High Performance Computing, Networking, Storage and Analysis*, page 87. ACM, November 2013.

[12] C. Karakoyunlu, D. Kimpe, P. Carns, K. Harms, R. Ross, and L. Ward. Towards a unified object storage foundation for scalable storage systems. In *Proceedings of the Fifth Workshop on Interfaces and Architectures for Scientific Data Storage (IASDS)*, September 2013.

[13] J. Soumagne, D. Kimpe, J. Zounmevo, M. Chaarawi, Q. Koziol, A. Afsahi, and R. Ross. Mercury: Enabling remote procedure call for high-performance computing. In *Proceedings of the IEEE Cluster Conference*, September 2013.

[14] J. A. Zounmevo, D. Kimpe, R. Ross, and A. Afsahi. Using MPI in high-performance computing services. In *Proceedings of the 20th European MPI Users' Group Meeting*, pages 43–48, September 2013.

[15] P. Carns, Y. Yao, K. Harms, R. Ross, and K. Antypas. Production I/O characterization on the Cray XE6. In *Proceedings of the 2013 Cray User Group Conference (CUG2013)*, May 2013.

[16] C. Sigovan, R. Ross, C. Muelder, K.-L. Ma, K. Iskra, and J. Cope. A visual network analysis method for large scale parallel I/O systems. In *Proceedings of the 27th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, May 2013.

[17] P. Carns, K. Harms, D. Kimpe, J. M. Wozniak, R. Ross, L. Ward, M. Curry, R. Klundt, G. Danielson, C. Karakoyunlu, et al. A case for optimistic coordination in hpc storage systems. In *7th Parallel Data Storage Workshop (PDSW 2012)*, November 2012.

[18] D. Goodell, S. J. Kim, R. Latham, M. Kandemir, and R. Ross. An evolutionary path to object storage access. In *Proceedings of the 7th Parallel Data Storage Workshop*, Salt Lake City, UT, November 2012.

[19] J. Jenkins, E. Schendel, S. Lakshminarasimhan, D. A. B. II, T. Rogers, S. Ethier, R. Ross, S. Klasky, and N. F. Samatova. Byte-precision level of detail processing for variable precision analytics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC12)*, Salt Lake City, UT, November 2012.

[20] S. Kumar, V. Vishwanath, P. Carns, J. A. Levine, R. Latham, G. Scorzelli, H. Kolla, R. Grout, R. Ross, M. E. Papka, J. Chen, and V. Pascucci. Efficient data restructuring and aggregation for I/O acceleration in PIDX. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC12)*, Salt Lake City, UT, November 2012.

[21] M. Mubarak, C. D. Carothers, R. B. Ross, and P. Carns. Modeling a million-node dragonfly network using massively parallel discrete event simulation. In *Proceedings of the 3rd International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS12) held as part of SC12*, November 2012.

[22] Z. Gong, T. Rogers, J. Jenkins, H. Kolla, S. Ethier, J. Chen, R. Ross, S. Klasky, and N. F. Samatova. MLOC: Multi-level layout optimization framework for compressed scientific data exploration with heterogeneous access patterns. In *Proceedings of the 41st International Conference on Parallel Processing (ICPP)*, pages 239–248, October 2012.

[23] T. Peterka and R. Ross. Versatile communication algorithms for data analysis. In *Special Session on Improving MPI User and Developer Interaction (IMUDI) at the 19th European MPI Users' Group Meeting*, Vienna, Austria, September 2012.

[24] T. Ilsche, J. Schuchart, J. Cope, D. Kimpe, T. Jones, A. Knüpfer, K. Iskra, R. Ross, W. Nagel, and S. Poole. Enabling event tracing at leadership-class scale through I/O forwarding middleware. In *Proceedings of the 21st international symposium on High-Performance Parallel and Distributed Computing*, pages 49–60, Delft, Netherlands, June 2012.

[25] D. Kimpe, P. Carns, K. Harms, J. M. Wozniak, S. Lang, and R. Ross. AESOP: Expressing concurrency in high-performance system software. In *Proceedings of the 7th International Conference on Networking, Architecture and Storage (NAS)*, pages 303–312, Fujian, China, June 2012.

[26] D. Kimpe, K. Mohror, A. Moody, B. van Essen, M. Gokhale, R. Ross, and B. R. de Supinski. Integrated in-system storage architecture for high performance computing. In *Proceedings of the 2nd International Workshop on Runtime and Operating Systems for Supercomputers*, Venice, Italy, June 2012.

[27] E. R. Schendel, S. V. Pendse, J. Jenkins, D. A. B. II, Z. Gong, S. Lakshminarasimhan, Q. Liu, H. Kolla, J. Chen, S. Klasky, R. Ross, and N. F. Samatova. ISOBAR hybrid compression-i/o interleaving for large-scale parallel i/o optimization. In *Proceedings of the 21st international symposium on High-Performance Parallel and Distributed Computing (HPDC)*, pages 61–72, Delft, Netherlands, June 2012.

[28] A. Gyulassy, V. Pascucci, T. Peterka, and R. Ross. The parallel computation of Morse-Smale complexes. In *Proceedings of the 26th International Parallel and Distributed Computing Symposium (IPDPS)*, pages 484–495, Boston, MA, May 2012.

[29] N. Liu, J. Cope, P. Carns, C. Carothers, R. Ross, G. Grider, A. Crume, and C. Maltzahn. On the role of burst buffers in leadership-class storage systems. In *Proceedings of the 2012 IEEE Conference on Massive Data Storage*, Pacific Grove, CA, April 2012.

[30] E. R. Schendel, Y. Jin, N. Shah, J. Chen, C. Chang, S.-H. Ku, S. Ethier, S. Klasky, R. Latham, R. Ross, and N. F. Samatova. ISOBAR preconditioner for effective and high-throughput lossless data compression. In *Proceedings of the 28th IEEE International Conference on Data Engineering (ICDE)*, Washington, DC, April 2012.

[31] Y. Jin, S. Lakshminarasimhan, N. Shah, Z. Gong, C. S. Chang, J. Chen, S. Ethier, H. Kolla, S.-H. Ku, S. Klasky, R. Latham, R. Ross, K. Schuchardt, and N. F. Samatova. S-preconditioner for multi-fold data reduction with guaranteed user-controlled accuracy. In *Proceedings of the 11th International Conference on Data Mining (ICDM)*, pages 290–299, Vancouver, Canada, December 2011.

[32] S. Lakshminarasimhan, J. Jenkins, I. Arkatkar, Z. Gong, H. Kolla, S.-H. Ku, S. Ethier, J. Chen, C. Chang, S. Klasky, R. Latham, R. Ross, and N. F. Samatova. ISABELA-QA: Query-driven analytics with ISABELA-compressed extreme-scale scientific data. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage, and Analysis (SC11)*, Seattle, WA, November 2011.

[33] M. Rodriguez, L. Ortiz, Y. Jia, K. Yoshii, R. Ross, and P. Beckman. Wireless sensor network for data center environmental monitoring. In *Proceedings of the Fifth International Conference on Sensing Technology (ICST)*, November 2011.

[34] W. Tantisiriroj, S. Patil, G. Gibson, S. W. Son, S. J. Lang, and R. B. Ross. On the duality of data-intensive file system design: Reconciling HDFS and PVFS. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC11)*, Seattle, WA, November 2011.

[35] N. Liu, C. Carothers, J. Cope, P. Carns, R. Ross, A. Crume, and C. Maltzahn. Modeling a leadership-scale storage system. In *Proceedings of the 9th International Conference on Parallel Processing and Applied Mathematics 2011 (PPAM 2011)*, October 2011.

[36] T. Peterka, R. Ross, W. Kendall, A. Gyulassy, V. Pascucci, H.-W. Shen, T.-Y. Lee, and A. Chaudhuri. Scalable parallel building blocks for custom data analysis. In *Proceedings of LDAV 2011*, Providence, RI, October 2011.

[37] S. Kumar, V. Vishwanath, P. Carns, B. Summa, G. Scorzelli, V. Pascucci, R. Ross, J. Chen, H. Kolla, and R. Grout. PIDX: Efficient parallel I/O for multi-resolution multi-dimensional scientific datasets. In *Proceedings of IEEE Cluster 2011*, Austin, TX, September 2011.

[38] B. Welton, J. Cope, D. Kimpe, C. Patrick, K. Iskra, and R. Ross. Improving I/O forwarding throughput with data compression. In *Proceedings of the Workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS) 2011*, Austin, TX, September 2011.

[39] S. Lakshminarasimhan, N. Shah, S. Ethier, S. Klasky, R. Latham, R. Ross, and N. F. Samatova. Compressing the incompressible with ISABELA:in-situ reduction of spatio-temporal data. In *Proceedings of EuroPar 2011*, Bordeaux, France, August/September 2011.

[40] C. Muelder, C. Sigovan, K.-L. Ma, J. Cope, S. Lang, P. B. Kamil Iskra, and R. Ross. Visual analysis of I/O system behavior for high end computing. In *Proceedings of the Workshop on Large-Scale System and Application Performance (LSAP 2011)*, June 2011.

[41] P. Carns, K. Harms, W. Allcock, C. Bacon, R. Latham, S. Lang, and R. Ross. Understanding and improving computational science storage access through continuous characterization. In *Proceedings of 27th IEEE Conference on Mass Storage Systems and Technologies (MSST 2011)*, May 2011.

[42] T. Peterka, R. Ross, B. Nouanesengsey, T.-Y. Lee, H.-W. Shen, W. Kendall, and J. Huang. A study of parallel particle tracing for steady-state and time-varying flow fields. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium*, Anchorage, AK, May 2011.

[43] S. W. Son, S. Lang, R. Latham, R. Ross, and R. Thakur. Reliable MPI-IO through layout-aware replication. In *Proceedings of the IEEE International Workshop on Storage Network Architecture and Parallel I/O*, May 2011.

[44] S. Kumar, V. Pascucci, V. Vishwanath, P. Carns, R. Latham, T. Peterka, M. Papka, and R. Ross. Towards parallel access of multi-dimensional, multiresolution scientific data. In *Proceedings of 2010 Petascale Data Storage Workshop*, November 2010.

[45] V. Vishwanath, M. Hereld, K. Iskra, D. Kimpe, V. Morozov, M. E. Papka, R. Ross, and K. Yoshii. Accelerating I/O forwarding in IBM Blue Gene/P systems. In *Proceedings of Supercomputing*, November 2010.

[46] P. Carns, R. Ross, and S. Lang. Object storage semantics for replicated concurrent-writer file systems. In *Proceedings of the Workshop on Interfaces and Abstractions for Scientific Data Storage*, September 2010.

[47] D. Kimpe, D. Goodell, and R. Ross. MPI datatype marshalling: A case study in datatype equivalence. In *Proceedings of EuroMPI*, September 2010.

[48] K. Ohta, D. Kimpe, J. Cope, K. Iskra, R. Ross, and Y. Ishikawa. Optimization techniques at the I/O forwarding layer. In *Proceedings of the IEEE International Conference on Cluster Computing*, September 2010.

[49] A. Shoshani, S. Klasky, and R. Ross. Scientific data management: Challenges and approaches in the extreme scale era. In *SciDAC 2010, Journal of Physics: Conference Series*, Chattanooga, TN, July 2010.

[50] J. Cope, K. Iskra, D. Kimpe, and R. Ross. Grids and HPC: Not as different as you might think? In *PARA 2010*, June 2010.

[51] J. Wozniak, S. W. Son, and R. Ross. Distributed object storage rebuild analysis via simulation with GOBS. In *Workshop on Fault-Tolerance for HPC at Extreme Scale*, June 2010.

[52] W. Kendall, T. Peterka, J. Huang, H.-W. Shen, and R. Ross. Accelerating and benchmarking radix-k image compositing at large scale. In *Proceedings of the Eurographics Symposium on Parallel Graphics and Visualization*, May 2010.

[53] S. W. Son, S. Lang, P. Carns, R. Ross, R. Thakur, B. Ozisikylimaz, P. Kumar, W.-K. Liao, and A. Choudhary. Enabling active storage on parallel I/O software stacks. In *Proceedings of the IEEE Symposium on Mass Storage Systems and Technologies*, May 2010.

[54] W. Kendall, M. Glatter, J. Huang, T. Peterka, R. Latham, and R. Ross. Terascale data organization for discovering multivariate climatic trends. In *Proceedings of Supercomputing*, November 2009.

[55] S. Lang, P. Carns, R. Latham, R. Ross, K. Harms, and W. Allcock. I/O performance challenges at leadership scale. In *Proceedings of Supercomputing*, November 2009.

[56] S. Narayan, J. Chandy, S. Lang, P. Carns, and R. Ross. Uncovering errors: The cost of detecting silent data corruption. In *Proceedings of the Petascale Data Storage Workshop*, November 2009.

[57] T. Peterka, D. Goodell, R. Ross, H.-W. Shen, and R. Thakur. A configurable algorithm for parallel image-compositing applications. In *Proceedings of Supercomputing*, November 2009.

[58] N. Ali, P. Carns, K. Iskra, D. Kimpe, S. Lang, R. Latham, and R. Ross. Scalable I/O forwarding framework for high-performance computing systems. In *IEEE International Conference on Cluster Computing (Cluster 2009)*, New Orleans, LA, September 2009.

[59] J. Blas, F. Isaila, J. Carretero, R. Latham, and R. Ross. Multiple-level MPI file write-back and prefetching for Blue Gene systems. In *Proc. of the 16th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2009)*, Espoo, Finland, September 2009.

[60] P. Carns, R. Latham, R. Ross, K. Iskra, S. Lang, and K. Riley. 24/7 characterization of petascale I/O workloads. In *Proceedings of the First Workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS)*, New Orleans, LA, September 2009.

[61] K. Gao, W. keng Liao, A. Choudhary, R. Ross, and R. Latham. Combining I/O operations for multiple array variables in parallel netCDF. In *Proceedings of 2009 Workshop on Interfaces and Architectures for Scientific Data Storage*, New Orleans, LA, September 2009.

[62] K. Gao, W. keng Liao, A. Nisar, A. Choudhary, R. Ross, and R. Latham. Using subfiling to improve programming flexibility and performance of parallel shared-file I/O. In *Proc. ICPP 09*, Vienna, Austria, September 2009.

[63] S. Lang, R. Latham, D. Kimpe, and R. Ross. Interfaces for coordinated access in the file system. In *Proceedings of 2009 Workshop on Interfaces and Architectures for Scientific Data Storage*, New Orleans, LA, September 2009.

[64] T. Peterka, H. Yu, R. Ross, K.-L. Ma, and R. Latham. End-to-end study of parallel volume rendering on the IBM Blue Gene/P. In *Proc. ICPP 09*, Vienna, Austria, September 2009.

[65] R. Ross, R. Latham, W. Gropp, E. Lusk, and R. Thakur. Processing MPI datatypes outside MPI. In *Proc. of the 16th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2009)*, Espoo, Finland, September 2009.

[66] A. Choudhary, W.-K. Liao, K. Gao, A. Nisar, R. Ross, R. Thakur, and R. Latham. Scalable I/O and analytics. In *SciDAC 2009, Journal of Physics: Conference Series*, San Diego, CA, July 2009.

[67] T. Peterka, R. Ross, H.-W. Shen, K.-L. Ma, W. Kendall, and H. Yu. Parallel visualization on leadership computing resources. In *SciDAC 2009, Journal of Physics: Conference Series*, San Diego, CA, July 2009.

[68] F. Isaila, J. G. Blas, J. Carretero, R. Latham, S. Lang, and R. Ross. Latency hiding file I/O for Blue Gene systems. In *Proceedings of the 9th IEEE International Symposium on Cluster Computing and the Grid*, May 2009.

[69] P. Carns, S. Lang, R. Ross, M. Vilayannur, J. Kunkel, and T. Ludwig. Small-file access in parallel file systems. In *Proceedings of the 23rd IEEE International Parallel and Distributed Processing Symposium*, April 2009.

[70] G. Grider, J. Nunez, J. Bent, S. Poole, R. Ross, and E. Felix. Coordinating government funding of file system and I/O research through the high end computing university research activity. In *SIGOPS Operating Systems Review*, January 2009.

[71] T. Peterka, R. Ross, H. Yu, K. Ma, R. Kooima, and J. Girado. Autostereoscopic display of large-scale scientific visualization. In *Proceedings of SPIE SD&A XX Conference*, San Jose, CA, January 2009.

[72] T. Peterka, R. Ross, H. Yu, K. Ma, W. Kendall, and J. Huang. Assessing improvements to the parallel volume rendering pipeline at large scale. In *Proceedings of Supercomputing 2008 Ultrascale Visualization Workshop*, Austin, TX, November 2008.

[73] W. Gropp, D. Kimpe, R. B. Ross, R. Thakur, and J. L. Träff. Self-consistent MPI-IO performance requirements and expectations. In *Proc. of the 15th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2008)*, September 2008.

[74] P. Gu, J. Wang, and R. Ross. Bridging the gap between parallel file systems and local file systems: A case study with PVFS. In *37th International Conference on Parallel Processing*, pages 554–561, September 2008.

[75] R. Ross, T. Peterka, H. Shen, Y. Hong, K. Ma, H. Yu, and K. Moreland. Parallel I/O and visualization at extreme scale. In *SciDAC 2008, Journal of Physics: Conference Series*, July 2008.

[76] T. Peterka, H. Yu, R. Ross, and K. Ma. Parallel volume rendering on the IBM Blue Gene/P. In *Proceedings of Eurographics Symposium on Parallel Graphics and Visualization 2008 (EGPGV'08)*, Crete, Greece, April 2008.

[77] A. Ching, W. Liao, A. Choudhary, R. Ross, and L. Ward. Noncontiguous locking techniques for parallel file systems. In *Proceedings of the 2007 ACM/IEEE conference on Supercomputing*, November 2007.

[78] D. Kimpe, R. Ross, S. Vandewalle, and S. Poedts. Transparent log-based data storage in MPI-IO applications. In *Proc. of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, September 2007.

[79] R. Latham, W. Gropp, R. B. Ross, and R. Thakur. Extending the MPI-2 generalized request interface. In *Proc. of the 14th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2007)*, pages 223–232, September 2007.

[80] K.-L. Ma, R. B. Ross, J. Huang, G. Humphreys, N. Max, K. Moreland, J. Owens, and H.-W. Shen. Ultra-scale visualization: Research and education. In *SciDAC 2007, Journal of Physics: Conference Series*, 2007.

[81] A. Shoshani, I. Altintas, A. Choudhary, T. Critchlow, C. Kamath, B. Ludascher, J. Nieplocha, S. Parker, R. B. Ross, N. Samatova, and M. Vouk. SDM center technologies for accelerating scientific discoveries. In *SciDAC 2007, Journal of Physics: Conference Series*, 2007.

[82] K. Coloma, A. Ching, A. Choudhary, W.-K. Liao, R. Ross, R. Thakur, and H. L. Ward. A new flexible MPI collective I/O implementation. In *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2006)*, September 2006.

[83] R. Latham, R. B. Ross, and R. Thakur. Can MPI be used for persistent parallel services? In *Proceedings of the 13th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2006)*, pages 275–284, September 2006.

[84] J. Lee, R. B. Ross, S. Atchley, M. Beck, , and R. Thakur. MPI-IO/L: Efficient remote I/O for MPI-IO via logistical networking. In *Proceedings of the 20th IEEE International Parallel and Distributed Processing Symposium (IPDPS 2006)*, Rhodes Island, Greece, April 2006.

[85] H. Yu, R. Sahoo, C. Howson, G. Almasi, J. Castanos, M. Gupta, J. Moreira, J. Parker, T. Engelsiepen, R. Ross, et al. High performance file I/O for the Blue Gene/L supercomputer. In *The Twelfth International Symposium on High-Performance Computer Architecture*, pages 187–196, 2006.

[86] R. Latham, R. Ross, and R. Thakur. Implementing MPI-IO shared file pointers without file system support. In *Proceedings of EuroPVM/MPI 2005*, September 2005.

[87] R. Thakur, R. Ross, and R. Latham. Implementing byte-range locks using MPI one-sided communication. In *Proceedings of the 12th European PVM/MPI Users' Group Meeting (Euro PVM/MPI 2005), Recent Advances in Parallel Virtual Machine and Message Passing Interface, Lecture Notes in Computer Science, LNCS 3666, Springer*, pages 119–128, September 2005.

[88] R. Ross, R. Latham, W. Gropp, R. Thakur, and B. Toonen. Implementing MPI-IO atomic mode without file system support. In *Proceedings of CCGrid 2005*, May 2005.

[89] P. H. Carns, W. B. Ligon III, R. B. Ross, and P. Wyckoff. BMI: A network abstraction layer for parallel I/O. In *Workshop on Communication Architecture for Clusters, Proceedings of IPDPS '05*, Denver, CO, April 2005.

[90] R. B. Ross, R. Thakur, and A. Choudhary. Achievements and challenges for I/O in computational science. In *SciDAC 2005: Scientific Discovery Through Advanced Computing, Journal of Physics: Conference Series*, pages 501–509, 2005.

[91] W. Gropp, R. B. Ross, , and N. Miller. Providing efficient I/O redundancy in MPI environments. In *Proceedings of EuroPVM/MPI 2004*, September 2004.

[92] R. Latham, R. Ross, and R. Thakur. The impact of file systems on MPI-IO scalability. In *Proceedings of EuroPVM/MPI 2004*, September 2004.

[93] J. Lee, X. Ma, R. B. Ross, R. Thakur, and M. Winslett. RFS: Efficient and flexible remote file access for MPI-IO. In *Proceedings of Cluster 2004*, September 2004.

[94] J. Wu, P. Wyckoff, D. Panda, and R. Ross. Unifier: Unifying cache management and communication buffer management for PVFS over InfiniBand. In *Proceedings of CCGrid2004*, Chicago, April 2004.

[95] M. Vilayannur, R. B. Ross, P. H. Carns, R. Thakur, and A. Sivasubramaniam. On the performance of the POSIX I/O interface to PVFS. In *12th Euromicro Conference on Parallel, Distributed and Network-Based Processing (PDP'04)*, pages 332–339, Coruna, Spain, February 2004.

[96] A. Ching, A. Choudhary, W. Liao, R. Ross, and W. Gropp. Efficient structured data access in parallel file systems. In *Proceedings of Cluster 2003*, Hong Kong, November 2003.

[97] J. Li, W. keng Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, , and M. Zingale. Parallel netCDF: A high-performance scientific I/O interface. In *Proceedings of SC2003*, November 2003.

[98] R. Ross, N. Miller, and W. Gropp. Implementing fast and reusable datatype processing. In *Proceedings of the 10th EuroPVM/MPI Conference*, September 2003.

[99] A. Ching, A. Choudhary, K. Coloma, W. keng Liao, R. Ross, and W. Gropp. Noncontiguous I/O accesses through MPI-IO. In *Proceedings of the Third IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGrid2003)*, May 2003.

[100] M. Vilayannur, A. Sivasubramaniam, M. Kandemir, R. Thakur, and R. Ross. Discretionary caching for I/O on clusters. In *Proceedings of the Third IEEE/ACM International Symposium on Cluster Computing and the Grid*, pages 96–103, Tokyo, Japan, May 2003. IEEE Computer Society Press.

[101] A. Ching, A. Choudhary, W. keng Liao, R. Ross, and W. Gropp. Noncontiguous I/O through PVFS. In *Proceedings of the 2002 IEEE International Conference on Cluster Computing*, September 2002.

[102] R. Ross, D. Nurmi, A. Cheng, and M. Zingale. A case study in application I/O on linux clusters. In *Proceedings of SC2001*, November 2001.

[103] P. H. Carns, W. B. Ligon III, R. B. Ross, and R. Thakur. PVFS: A parallel file system for Linux clusters. In *Proceedings of the 4th Annual Linux Showcase and Conference*, pages 317–327, Atlanta, GA, October 2000. USENIX Association.

[104] W. B. Ligon III and R. B. Ross. An overview of the parallel virtual file system. In *Proceedings of the 1999 Extreme Linux Workshop*, Monterey, CA, June 1999.

[105] P. H. Carns, W. B. Ligon III, S. McMillan, and R. B. Ross. An evaluation of message passing implementations on Beowulf workstations. In *Proceedings of the IEEE Aerospace Conference*, Snowmass, CO, March 1999.

[106] M. Cettei, W. Ligon, and R. Ross. Support for parallel out of core applications on Beowulf workstations. In *Proceedings of the 1998 IEEE Aerospace Conference*, March 1998.

[107] R. Geist and R. Ross. Disk scheduling revisited: Can $O(N^2)$ algorithms compete? In *Proceedings of the 35th Annual ACM Southeast Conference*, April 1997.

[108] W. B. Ligon III and R. B. Ross. Implementation and performance of a parallel file system for high performance distributed applications. In *Proceedings of the Fifth IEEE International Symposium on High Performance Distributed Computing (HPDC)*, Syracuse, NY, August 1996.

## Invited Talks at Major Conferences

[1] R. B. Ross. Storage in an exascale world. Presented at the IEEE International Workshop on Storage Network Architecture and Parallel I/Os (SNAPI), Incline Village, NV, May 2010.

[2] R. B. Ross. Visualization and parallel I/O at extreme scale. Presented at the 2008 SciDAC Conference, Seattle, WA, July 2008.

[3] R. B. Ross. PVFS: The parallel virtual file system. Presented at the Storage Networking Industry Association Developer Solutions Conference & Showcase, San Jose, CA, August 2005.

[4] R. B. Ross. Achievements and challenges for I/O in computational science. Presented at the 2005 SciDAC Conference, San Francisco, CA, June 2005.

[5] R. B. Ross. The parallel I/O software crisis. Presented at ISC, Heidelberg, Germany, June 2005.

[6] R. B. Ross. PVFS2: Parallel I/O for scientific applications. Presented at ClusterWorld 2004, San Jose, CA, 2004.

[7] R. B. Ross. Providing parallel I/O on linux clusters. Presented at the Second Annual Linux Storage Management Workshop, Miami, FL, October 2000.

## Tutorials

[1] R. Latham, R. Ross, and Q. Koziol. HPC I/O for computational scientists. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2014.

[2] R. Latham, K. Antypas, R. Ross, and B. Welch. Parallel I/O in practice. SC 2013, Denver, CO, November 2013.

[3] R. Latham, R. Ross, and Q. Koziol. HPC I/O for computational scientists. Argonne Training Program on Extreme-Scale Computing (ATPESC), St. Charles, IL, August 2013.

[4] R. Ross. (Big) data in computational science. Short Course on The Materials Genome, Current Practice and Future Promise, NSF Summer Institute on Nanomechanics, Nanomaterials, and Micro/Nanomanufacturing, Evanston, IL, June 2013.

[5] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2012, Salt Lake City, UT, November 2012.

[6] R. Latham, K. Antypas, R. Ross, and B. Welch. Parallel I/O in practice. SC 2012, Salt Lake City, UT, November 2012.

[7] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2011, Seattle, WA, November 2011.

[8] R. Latham, K. Antypas, R. Ross, and B. Welch. Parallel I/O in practice. SC 2011, Seattle, WA, November 2011.

[9] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2010, New Orleans, LA, November 2010.

[10] R. Latham, R. Ross, M. Unangst, and B. Welch. Parallel I/O in practice. SC 2010, New Orleans, LA, November 2010.

[11] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2009, Portland, OR, November 2009.

[12] R. Latham, R. Ross, M. Unangst, and B. Welch. Parallel I/O in practice. SC 2009, Portland, OR, November 2009.

[13] R. Latham and R. Ross. Parallel I/O in practice. SciDAC Tutorials Day, San Diego, CA, June 2009.

[14] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2008, Austin, TX, November 2008.

[15] R. Latham, R. Ross, M. Unangst, and B. Welch. Parallel I/O in practice. SC 2008, Austin, TX, November 2008.

[16] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI. SC2007, Reno, NV, November 2007.

[17] R. Latham, W. Loewe, R. Ross, and R. Thakur. Parallel I/O in practice. SC2007, Reno, NV, November 2007.

[18] R. Latham and R. Ross. Parallel I/O: Not your job. CScADS Workshop on Petascale Architectures and Performance Strategies, Snowbird, UT, July 2007.

[19] R. Latham and R. Ross. Parallel I/O in practice. SciDAC 2007 Tutorials Workshop, Boston, MA, June 2007.

[20] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI: I/O and one-sided communication. SC2006, Tampa, FL, November 2006.

[21] R. Latham, W. Loewe, R. Ross, and R. Thakur. Parallel I/O in practice. SC2006, Tampa, FL, November 2006.

[22] R. Latham and R. Ross. Parallel I/O in practice. Cluster 2006, Barcelona, Spain, September 2006.

[23] R. Ross and J. Worringen. High-performance parallel I/O. EuroPVM/MPI 2006, Bonn, Germany, September 2006.

[24] R. Ross. Parallel I/O in practice. Sandia National Laboratories, Albuquerque, NM, July 2006.

[25] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI: I/O and one-sided communication. SC2005, Seattle, WA, November 2005.

[26] R. Latham, W. Loewe, R. Ross, and R. Thakur. Parallel I/O in practice. SC2005, Seattle, WA, November 2005.

[27] R. Latham and R. Ross. High-performance I/O for scientific applications. CCGrid 2005, Cardiff, UK, May 2005.

[28] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Advanced MPI: I/O and one-sided communication. SC2004, Pittsburgh, PA, November 2004.

[29] R. B. Ross. High-performance I/O for scientific applications. ClusterWorld 2004, San Jose, CA, April 2004.

[30] R. B. Ross and R. Thakur. Using MPI-2: A tutorial on advanced features of the message-passing interface standard. CCGrid 2004, Chicago, April 2004.

[31] R. Ross and R. Thakur. Using MPI-2: A tutorial on advanced features of the message-passing interface standard. Grid and Cluster Computing Conference (GCC) 2003, Shanghai, China, December 2003.

[32] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Using MPI-2: A tutorial on advanced features of the message-passing interface standard. SC2003, Phoenix, AZ, November 2003.

[33] W. Gropp, E. Lusk, R. Ross, and R. Thakur. High-level programming with MPI. EuroPVM/MPI 2003, Venice, September 2003.

[34] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Using MPI-2: A tutorial on advanced features of the message-passing interface standard. SC2002, Baltimore, MD, November 2002.

[35] W. Gropp, E. Lusk, R. Ross, and R. Thakur. Using MPI-2: A tutorial on advanced features of the message-passing interface standard. SC2001, Denver, CO, November 2001.

[36] W. B. L. III and R. B. Ross. The parallel virtual file system for commodity clusters. IEEE Cluster 2001, Newport Beach, CA, October 2001.

[37] R. Pennington, P. Kovatch, B. Maccabe, D. Bader, and R. Ross. Design and analysis of high performance clusters. SC2000, Dallas, TX, November 2000.

## Seminars

[1] R. B. Ross. Big Data and DOE science. Presented at INRIA, Rennes, France, June 2014.

[2] R. B. Ross. I/O workloads... Presented at the 11th Workshop of the INRIA-Illinois-ANL Joint Laboratory on Petascale Computing, Sophia Antipolis, France, June 2014.

[3] R. B. Ross. Future HPC systems and some implications for storage software. Presented at the 2014 Lustre Users Group Meeting, Panel Session, Miami, FL, USA, April 2014.

[4] R. B. Ross. HPC storage and data: Current state and future directions. Presented at the 2014 High-Performance Computing and Data-Intensive Geospatial Analytics Workshop, Argonne, IL, April 2014.

[5] R. B. Ross and D. Kimpe. Argonne data/storage activities. Presented at Western Digital, Irvine, CA, April 2014.

[6] R. B. Ross. High performance computing I/O systems: Overview and recent developments. Presented at Illinois Institute of Technology, Chicago, IL, USA, November 2013.

[7] R. B. Ross. Thinking past POSIX: persistent storage in extreme scale systems. Presented at the INRIA-Illinois Joint Lab Workshop, Urbana, IL, November 2013.

[8] R. B. Ross. Exascale storage and I/O: Current activities and possible directions. Presented at the First International Workshop on Strategic Development of High Performance Computers, Tsukuba, Japan, March 2013.

[9] R. B. Ross. Trends in HPC I/O and file systems. Presented at the INRIA-Illinois Joint Lab Workshop, Argonne, IL, November 2012.

[10] R. B. Ross. Computational science and big data: Connections and opportunities. Presented at the DOE Big Data Technology Summit, Washington, DC, October 2012.

[11] R. B. Ross. The DOE SciDAC institute for scalable data management, analysis, and visualization. Presented at the 7th Meeting of the Board on Research Data and Integration (BRDI), Washington, DC, August 2012.

[12] R. B. Ross. Big data and scientific computing: Some initial thoughts. Presented at the Ninth Workshop of the INRIA-Illinois Joint Laboratory on Petascale Computing, Lyon, France, June 2012.

[13] R. B. Ross. Scientific data management and analysis challenges. Presented at the Transforming Geant4 for the Future Workshop, Rockville, MD, May 2012.

[14] R. B. Ross. Research strategies for I/O scalability. Presented at the April 2012 Exascale Research Conference, Arlington, VA, April 2012.

[15] R. B. Ross. Resilience, storage, and exascale computing. Presented at the Inter-Agency Workshop on HPC Resilience at Extreme Scale, Catonsville, MD, February 2012.

[16] R. B. Ross. Fear and loathing in data storage. Presented during the panel on Scientific Data on the Path to Exascale at SC 2011, Seattle, WA, November 2011.

[17] R. B. Ross. Overcoming roadblocks to exascale storage. Presented during the Open Source File System BoF at SC 2011, Seattle, WA, November 2011.

[18] R. B. Ross. Storage architectures and abstractions for exascale systems. Presented at the INRIA-Illinois Petascale Computing Joint Lab Workshop, Urbana, IL, November 2011.

[19] R. B. Ross. Storage for extreme scale computing. Presented as part of the CS Seminar Series at Illinois Institute of Technology, Chicago, IL, November 2011.

[20] R. B. Ross. The exascale software center: Data storage and analysis. Presented at the Sixth Workshop of the International Exascale Software Project, San Francisco, CA, April 2011.

[21] R. B. Ross. Open source I/O software for HPC: A quick tour. Presented at the DOE Workshop for Industry Software Developers, Chicago, IL, March 2011.

[22] R. B. Ross. Planning for the Exascale Software Center. Presented at the DOE Advanced Scientific Computing Advisory Committee (ASCAC) meeting, Argonne, IL, November 2010.

[23] R. B. Ross. Preparing for exascale: Understanding HPC storage systems. Presented at the Workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS), Heraklion, Crete, Greece, September 2010.

[24] R. B. Ross. Data models and data analysis at exascale. Presented at the High-End Computing File Systems and I/O Conference, Arlington, VA, August 2010.

[25] R. B. Ross. Scientific computing at extreme scale. Presented at the University of Connecticut, Storrs, CT, June 2010.

[26] R. B. Ross. Applications, data, and the future of storage in computational science. Presented at the SCI Institute, University of Utah, Salt Lake City, UT, May 2010.

[27] R. B. Ross. Input/output (I/O) in computational science. Presented at the Computation Institute at the University of Chicago, Chicago, IL, February 2010.

[28] R. B. Ross. Extreme scale I/O systems. Presented at the IEEE Nuclear Science Symposium Data-Intensive Workshop, Orlando, FL, October 2009.

[29] R. B. Ross. The Scientific Data Management Center. Presented at the High-End Computing File Systems and I/O Conference, Arlington, VA, August 2009.

[30] R. B. Ross. Parallel I/O in practice. Presented at the CScADS Workshop on Leadership-class Machines, Parallel Applications, and Performance Strategies, Tahoe City, CA, July 2009.

[31] R. B. Ross. Meeting the needs of computational science at extreme scale. Presented at the University of Chicago at Santa Cruz, Santa Cruz, CA, January 2009.

[32] R. B. Ross. Parallel I/O and computational science at the largest scales. Presented at the Blue Waters Workshop, University of Illinois at Urbana-Champaign, Urbana, IL, October 2008.

[33] R. B. Ross. The reality of storage in computational science. Presented at Carnegie Mellon University, Pittsburgh, PA, April 2008.

[34] R. B. Ross. The SciDAC SDM center: Moving research into production. Presented at the High-End Computing File Systems and I/O Conference, Arlington, VA, August 2007.

[35] R. B. Ross. Pushing research into reality. Presented at the ACS Workshop, Linthicum, MD, June 2007.

[36] R. B. Ross. Storage at scale: Parallel I/O. Presented at the Electronic Visualization Laboratory at the University of Illinois at Chicago, Chicago, IL, May 2007.

[37] R. B. Ross. Parallel programming and MPI. Presented at Carnegie-Mellon University, Pittsburgh, PA, April 2007.

[38] R. B. Ross. Approaching petascale I/O. Presented at the Geosciences Application Requirements for Petascale Architectures (GARPA-2), San Diego, CA, February 2007.

[39] R. B. Ross. PVFS in production. Presented at Lawrence Livermore National Laboratory, Livermore, CA, February 2007.

[40] R. B. Ross. I/O at petascale: Enabling and understanding. Presented at Texas A&M University, College Station, TX, January 2007.

[41] R. B. Ross. A Department of Energy perspective on parallel I/O. Presented at The 2006 Workshop on Cluster Storage Technology (CluStor 2006), Heidelberg, Germany, September 2006.

[42] R. B. Ross. PVFS in production. Presented at Sandia National Laboratories, Albuquerque, NM, July 2006.

[43] R. B. Ross. Trends and techniques in parallel I/O systems. Presented at Invited Speaker Series, School of Computing and Information Sciences, Florida International University, Miami, FL, January 2006.

[44] R. B. Ross, T. Baer, A. Ching, D. Hildebrand, and R. Latham. PVFS2 birds of a feather session. Presented at SC2005, Seattle, WA, November 2005.

[45] R. B. Ross. Building effective I/O solutions for HPC. Presented at Computer Sciences and Mathematics Division Seminar, Oak Ridge National Laboratory, Oak Ridge, TN, August 2005.

[46] R. B. Ross. Building parallel file systems for computational science. Presented at Fulton HPC Distinguished Lecture Series, Arizona State University, Tempe, AZ, April 2005.

[47] R. B. Ross. PVFS2 and parallel I/O on BG/L. Presented at BG/L Consortium System Software Workshop, Salt Lake City, UT, February 2005.

[48] R. B. Ross. The future of parallel I/O systems in computational science. Presented at the Workshop on System-Integrated Load and Resource Management, University of Heidelberg, Heidelberg, Germany, November 2004.

[49] R. B. Ross, R. Latham, W. Ligon, and N. Miller. PVFS2 birds of a feather session. Presented at SC2004, Pittsburgh, PA, November 2004.

[50] R. B. Ross. Bridging the I/O gap: Matching I/O systems to application domains. Presented at Center for Computing Sciences Colloquium Series, Institute for Defense Analysis, Bowie, MD, March 2004.

[51] R. B. Ross. Connecting HPIO capabilities with domain specific needs. Presented at DOE Office of Science Data-Management Workshop, Menlo Park, CA, March 2004.

[52] R. B. Ross, W. Ligon, P. Carns, R. Latham, and N. Miller. PVFS birds of a feather session. Presented at SC2003, Phoenix, AZ, November 2003.

[53] R. B. Ross. Parallel I/O systems: Architecture and performance. Presented at Systems Seminar Series, Ohio State University, Columbus, OH, April 2003.

[54] R. B. Ross and W. B. Ligon. PVFS birds of a feather session. Presented at SC2002, Baltimore, MD, November 2002.

[55] R. B. Ross. Making best use of PVFS. Presented at Cluster Focus Group Meeting, Ohio Supercomputer Center, Columbus, OH, April 2002.

[56] R. B. Ross. Revisiting the parallel I/O problem. Presented at Lawrence Livermore National Laboratory, Livermore, CA, October 2001.

[57] R. B. Ross. Using parallel I/O on linux clusters. Presented at NASA Jet Propulsion Laboratory, High Performance Computing Group, Pasadena, CA, July 2001.

[58] R. B. Ross. Playing with parallel I/O on Linux clusters. Presented at University of Chicago, Chicago, IL, February 2001.

[59] R. B. Ross. Reactive scheduling for parallel I/O systems. Presented at Argonne National Laboratory, Chicago, IL, May 2000.

[60] R. B. Ross. Message passing and parallel file systems for Beowulf machines. Presented at NASA Goddard Space Flight Center, Greenbelt, MD, April 1999.

[61] R. B. Ross. The Parallel Virtual File System: Past, present, and future. Presented at Argonne National Laboratory, Chicago, IL, April 1999.

## Magazine Articles

[1] P. C. Wong, H.-W. Shen, C. R. Johnson, C. Chen, and R. B.Ross. The top 10 challenges in extreme-scale visual analytics. *Computer Graphics and Applications*, 32(4):63–67, 2012.

[2] W. Kendall, J. Huang, T. Peterka, R. Latham, and R. Ross. Toward a general I/O layer for parallel-visualization applications. *IEEE Computer Graphics and Applications*, 31(6), November/December 2011.

[3] J. Ahrens, B. Hendrickson, G. Long, S. Miller, R. Ross, and D. Williams. Data intensive science in the Department of Energy: Case studies and future challenges. *IEEE Computing in Science and Engineering*, 13(6):14–23, 2011.

[4] K.-L. Ma, C. Wang, H. Yu, K. Moreland, J. Huang, and R. Ross. Next-generation visualization technologies: Enabling discoveries at extreme scale. *SciDAC Review*, Spring 2009.

[5] N. Miller, R. Latham, R. B. Ross, and P. Carns. Improving cluster performance with PVFS2. *ClusterWorld Magazine*, 2(4), April 2004.

[6] R. Latham, N. Miller, R. B. Ross, and P. Carns. A next-generation parallel file system for Linux clusters. *LinuxWorld Magazine*, 2(1), January 2004.

[7] D. Becker, W. B. Ligon III, P. Merkey, and R. B. Ross. Beowulf: Low-cost supercomputing using Linux. *IEEE Software*, January/February 1999.

## Technical Reports and Whitepapers

[1] R. Ross, G. Grider, E. Felix, M. Gary, S. Klasky, R. Oldfield, G. Shipman, and J. Wu. Storage systems and input/output to support extreme scale science. Technical report, May 2015.

[2] S. Habib, R. Roser, C. Tull, B. Hendrickson, R. Ross, and A. Shoshani. Report on HEP/ASCR data summit. Technical report, April 2013.

[3] J. Ahrens, B. Hendrickson, G. Long, S. Miller, R. Ross, and D. Williams. Data intensive science in the department of energy. Technical report, October 2010.

[4] M. Vilayannur, S. Lang, R. Ross, R. Klundt, and L. Ward. Extending the POSIX I/O interface: A parallel file system perspective. Technical Report ANL/MCS-TM-302, Argonne National Laboratory, 2008.

[5] C. Johnson, R. Ross, S. Ahern, J. Ahrens, W. Bethel, K. L. Ma, M. Papka, J. van Rosendale, H. W. Shen, and J. Thomas. Visualization and knowledge discovery: Report from the DOE/ASCR workshop on visual analysis and data exploration at extreme scale. Technical report, Department of Energy Office of Advanced Scientific Computing Research, October 2007.